COMP4970/7970/7976 - Deep Learning

Spring 2018, MWF 9-9:50am, Shelby 1122

Instructor
Dr. Anh Nguyen
Email: anhnguyen@auburn.edu
Phone: +1.608.334.9754
Office Hours: MW: 10:00-11:00am Room: Shelby 3101R

Teaching Assistant
Chengfei Wang
Email: czw0078@auburn.edu
Office Hours: TR: 10:00-11:30am Room: Shelby 2105

Deep learning, has become ubiquitous in our society, with many real-world applications in search, image understanding, speech, medicine, or self-driving cars. This is an advanced machine learning course that will expose students to cutting-edge research in deep learning. Students will learn starting from the basics of neural networks (implementing and training their own neural networks via backpropagation e.g. just using numpy), to recent developments in the field. We will mainly discuss convolutional and recurrent neural networks and their applications in computer vision, generative modeling, and natural language processing. The course will be project-based, demanding in terms of programming, with emphasis on paper reading and presentations. Students will be given opportunities to work on exciting, publishable research projects.

Prerequisites: calculus, linear algebra, probability and statistics, intro to machine learning or equivalent, and programming skills in Python.

Objectives

1. Understand basic neural network architectures and implement them from scratch
2. Understand some of the latest advances in deep learning / neural networks.
3. Understand pros & cons of current deep learning research and apply deep learning to a problem of interests.

Textbooks and Notes

There are no required textbooks for this class. Optional reading materials are below:

- Book: Goodfellow, Bengio, and Courville, **Deep Learning**, MIT Press, 2017.
- Lecture notes: http://cs231n.stanford.edu

Grades

Your grade will be computed as follows:

4970 – **Undergrad** level:

1. Homework 1: 10%
2. Homework 2: 10%
3. Homework 3: 10%
4. Project: 40%
5. Presentation Questions & Reviews: 30%

7970/7976 – **Graduate** level:

1. Homework 1: 10%
2. Homework 2: 10%
3. Homework 3: 10%
4. Project: 40%
5. Paper Presentations: 20%
6. Presentation Reviews: 10%

Academic Integrity
All tests, homework assignments, and especially programming assignments are expected to be exclusively the work of the one student submitting the assignment. Now, let me say that again, in different language. Do not give anyone a line of your code. Do not copy a single line of code from any other student, or any web site, or any previous student. If you do, it is cheating. You could fail the course or suffer other penalties. If you have any doubt about whether what you are doing is ethical, ask me -- I won't penalize anyone for asking, or for following my advice. Be aware I sometimes use a program that compares all student programs for algorithmic (not textual) similarity.

Special Accommodations
Students who need special accommodations should make an appointment to discuss your needs during office hours as soon as possible. If you do not have an Accommodation Memo, but need special accommodations, contact the appropriate university office.

Tentative Schedule (next page)

Module	Week	Event	Date	Description	Presenter	Notes
	1	L1	1/10	Course introduction & logistics		
		L2	1/12	Image Classification		A1 out
Training CNNs	2	*No class*	1/15	*M.L. King Day*		
		L3	1/17	Loss Functions and Optimization		
		L4	1/19	Introduction to Neural Networks		
	3	L5	1/22	Convolutional Neural Networks		
		L6	1/24	Training Neural Networks, part I		
		L7	1/26	Training Neural Networks, part II		**A1 due**
	4	L8	1/29	Training Neural Networks, part III		A2 out
		L9	1/31	CNN Architectures		**Proposal due**
		Presentation	2/2	Stochastic Depth Network	Group 1	
Understanding CNNs	5	L10	2/5	Introduction to Caffe	Chengfei	
		L11	2/7	Visualizing and Understanding CNNs, part I		
		Presentation	2/9	Network Dissection (Zhou et al. 2017)	Group 2	
	6	L12	2/12	Visualizing and Understanding CNNs, part II		
		L13	2/14	Fooling CNNs		**A2 due**
		Presentation	2/16	Celebrity glasses CMU (Sharif et al. 2017)	Group 3	
AI for Creativity	7	*No class*	2/19	*Presidents' Day*		
		L14	2/21	DeepDream / Style Transfer / Smile Vector / Creative AI		A3 out
		L15	2/23	Introduction to Tensorflow & Keras	Chengfei	
Detection & Segmentation	8	L16	2/26	Detection CNNs		
		L17	2/28	Segmentation CNNs		
		Presentation	3/2	DeepMask	Group 4	
Recurrent Neural Networks	9	L18	3/5	Recurrent Neural Networks part I		
		L19	3/7	Recurrent Neural Networks part II		
		Presentation	3/9	Visualizing & Understanding RNNs (Karpathy et al. 2015)	Group 5	
		No class	3/12 3/14 3/16	*Spring Break*		
Sequence to sequence	10	L20	3/19	Machine Translation		
		L21	3/21	Speech		**A3 due**
		Presentation	3/23	Neural Programmer-Interpreter (Reed & Freitas, 2015)	Group 6	
Theory of Deep Learning	11	L22	3/26	Theory of Deep Learning I		
		L23	3/28	Theory of Deep Learning II		
		Presentation	3/30	Information Bottleneck (Tishby 2017)	Group 7	
Generative Models	12	L24	4/2	Intro to Generative Models / Unsupervised Learning		
		L25	4/4	Autoregressive / PixelCNN		
		Presentation	4/6	SketchRNN (Ha & Douglas. 2017)	Group 8	
	13	L26	4/9	Variational methods / VAE		
		L27	4/11	Generative Adversarial Networks (GAN)		
		Presentation	4/13	VQ-VAE (Oord et al. 2017)	Group 9	
Deep RL	14	L28	4/16	Deep Reinforcement Learning I		
		L29	4/18	Deep Reinforcement Learning II		
		Presentation	4/20	AlphaGo / AlphaZero (DeepMind 2017)	Group 10	
Project	15	Presentation	4/23	Project Day 1		**Project due**
		Presentation	4/25	Project Day 2		
		Presentation	4/27	Project Day 3		
	16			Final exam		